

ODA @ ENI Green Data Center

ODA: goals and requirements (year is 2008)

Facility analysis tool

Failure prevention

Energy-wise optimization

Fault detection

Collect any available data

High polling frequency

Long term online data retention



ODA: plate data

Up & running since facility building

350 multimeters, 1800 probes, 112 UPS, 4000 switches...

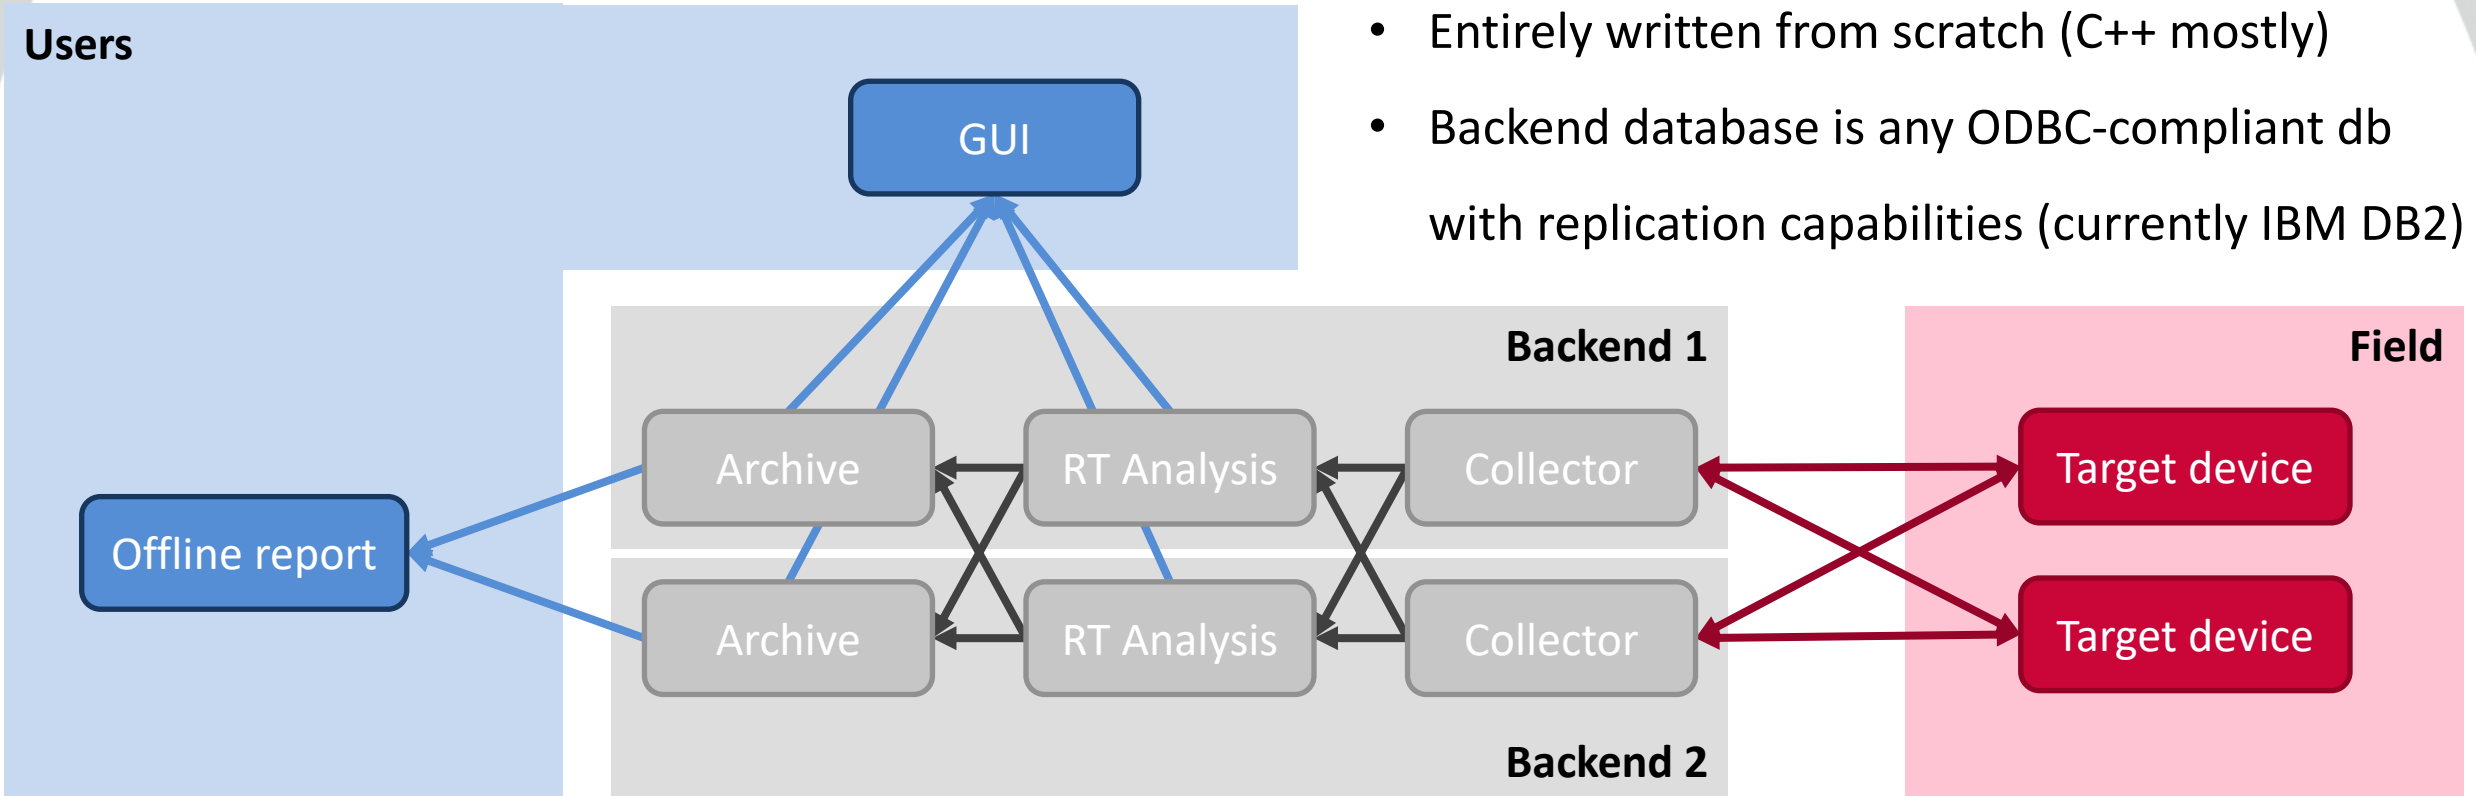
Since 2 years HPC hardware metrics added
(CPU, GPU, IB, storage...)

Currently >200k samples collected every 10 seconds (avg)

8 years online: 1.3×10^{12} samples, 1.3TB



ODA: architecture overview



- Entirely written from scratch (C++ mostly)
- Backend database is any ODBC-compliant db with replication capabilities (currently IBM DB2)

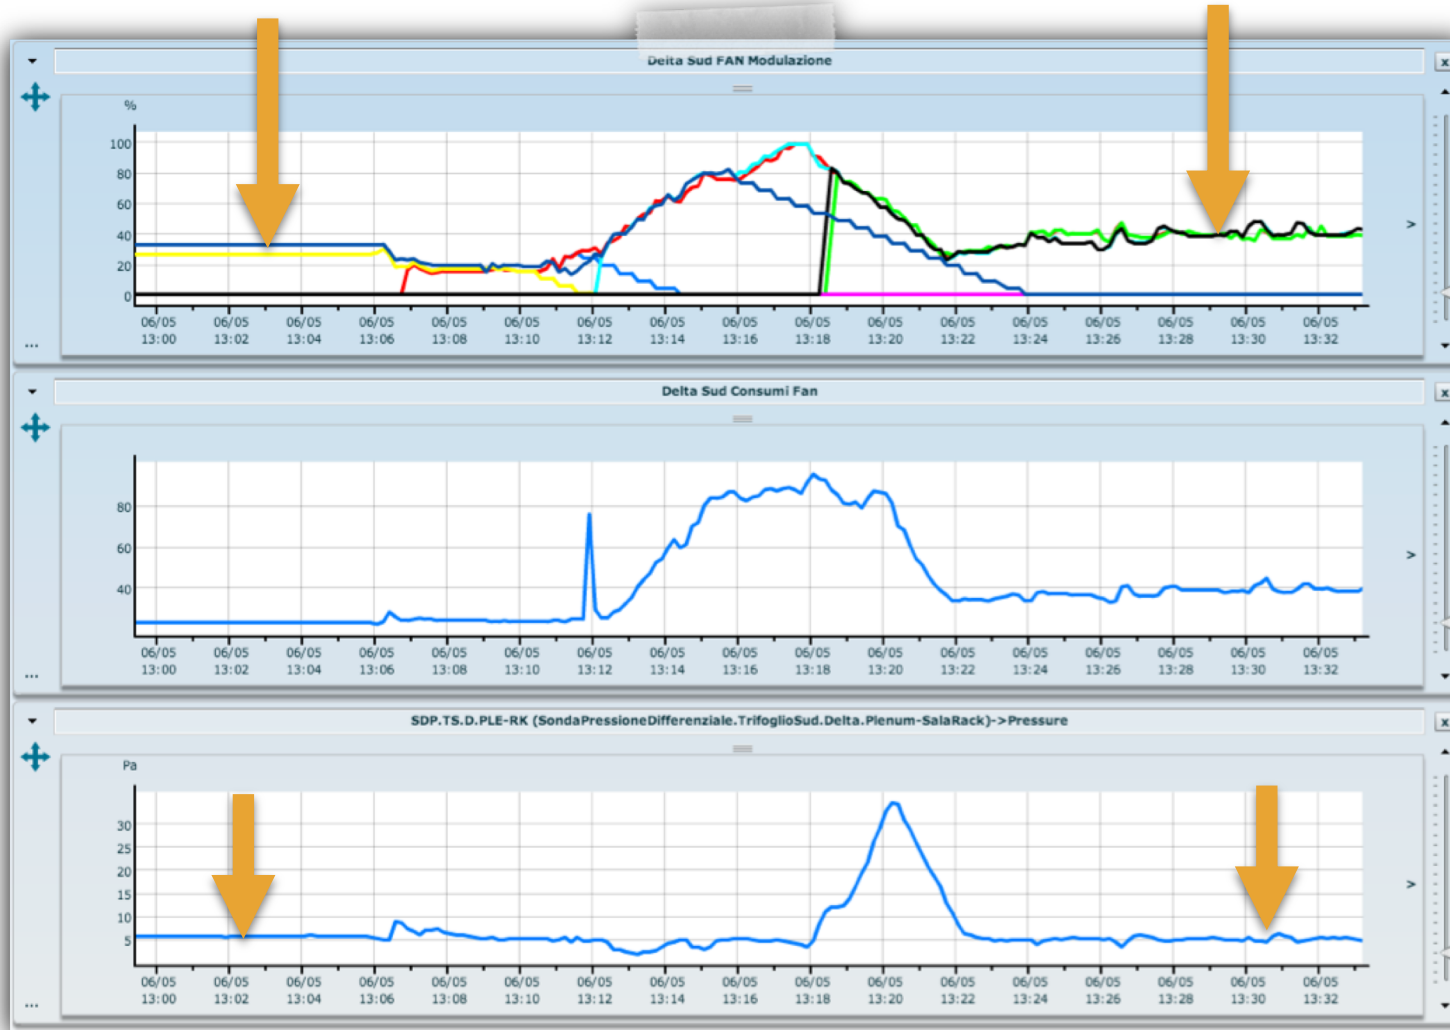
- Symmetric asynchronous data replication over the whole pipeline
- Allows online software fixes and upgrades (7 years service uptime)
- High efficiency – every backend runs on a single server

ODA: data life cycle

Step	Technology	Activity	Samples per record	Record count	Notes
Data collector	Custom	Data production	1	244k	
RT analysis	Custom	Alarm generation	1	244k	
Daily storage	Custom	Storage	~360	87k in-memory, 200k per table on disk	Data are scattered across multiple tables
Weekly storage	RDBMS	Storage	~7k	88k per table	Data arrays are compressed (lossless)
Monthly storage (up to 6 months)	RDBMS	Storage	10k	3m per table	
Long term	RDBMS	Storage	10k	<10m per table	Read-only tablespace



Use case: «hidden» turbines fault



Average load
2x26% -> 4x38%

Power drain
22Kw -> 38Kw

Pressure unchanged